

## Welcome United States Patent and Trademark Office

**#**■#Search Results

BROWSE

**SEARCH** 

**IEEE XPLORE GUIDE** 

-Search

SUPPORT

e-mail printer friendly

Results for "((attribute cache)<in>metadata)"

Your search matched 1 of 1325881 documents.

A maximum of 100 results are displayed, 25 to a page, sorted by Relevance in Descending order.

» Search Options

View Session History

New Search

» Key

IEEE JNL

IEEE Journal or Magazine

IEE JNL

IEE Journal or Magazine

IEEE CNF

IEEE Conference Proceeding

**IEE CNF** IEEE STD IEE Conference Proceeding

IEEE Standard

**Modify Search** 

((attribute cache)<in>metadata)

Check to search only within this results set

Display Format:

Citation

⑥ Citation & Abstract

view selected items

Select All Deselect All

1. Collective communication and communicators in mpi++

Kafura, D.; Huang, L.;

MPI Developer's Conference, 1996. Proceedings., Second

1-2 July 1996 Page(s):79 - 86

Digital Object Identifier 10.1109/MPIDC.1996.534097

AbstractPlus | Full Text: PDF(676 KB) | IEEE CNF

Rights and Permissions

Contact Us Privacy & Security IEEE.org

© Copyright 2006 IEEE - All Rights Reserved

Indexed by #Inspec

SUPPORT

**IEEE XPLORE GUIDE** 



## Welcome United States Patent and Trademark Office

SEARCH

**BROWSE** 

□□□Search Results Results for "((attribute cache <or> attribute caching)<in>metadata)" e-mail printer friendly Your search matched 1 of 1325881 documents. A maximum of 100 results are displayed, 25 to a page, sorted by Relevance in Descending order. » Search Options View Session History **Modify Search** ((attribute cache <or> attribute caching)<in>metadata) New Search "Search. Check to search only within this results set » Key Display Format: Citation Citation & Abstract IEEE JNL IEEE Journal or Magazine IEE JNL IEE Journal or Magazine view selected items Select All Deselect All IEEE CNF IEEE Conference Proceeding **IEE CNF** IEE Conference Proceeding 1. Collective communication and communicators in mpi++ Kafura, D.; Huang, L.; IEEE STD IEEE Standard MPI Developer's Conference, 1996. Proceedings., Second 1-2 July 1996 Page(s):79 - 86 Digital Object Identifier 10.1109/MPIDC.1996.534097 AbstractPlus | Full Text: PDF(676 KB) IEEE CNF Rights and Permissions

Indexed by #Inspec Help Contact Us Privacy & Security IEEE.org © Copyright 2006 IEEE - All Rights Reserved



## Welcome United States Patent and Trademark Office

**BROWSE SEARCH** IEEE XPLORE GUIDE SUPPORT **□□**Search Results Results for "(((attribute cache <or> attribute caching) <and> tib) <in>metadata)" e-mail Aprinter friendby Your search matched 0 documents. A maximum of 100 results are displayed, 25 to a page, sorted by Relevance in Descending order. » Search Options View Session History **Modify Search** (((attribute cache <or> attribute caching) <and> tlb) <in>metadata) New Search Search > Check to search only within this results set » Key Citation Citation & Abstract Display Format: **IEEE JNL** IEEE Journal or Magazine IEE JNL IEE Journal or Magazine **IEEE CNF** IEEE Conference Proceeding No results were found. **IEE CNF** IEE Conference Proceeding Please edit your search criteria and try again. Refer to the Help pages if you need assistance revising your search. **IEEE STD** IEEE Standard

Indexed by

Help Contact Us Privacy & Security IEEE.org

© Copyright 2006 IEEE - All Rights Reserved

Subscribe (Full Service) Register (Limited Service, Free) Login

Search: 

The ACM Digital Library O The Guide

"attribute cache" and TLB

SEARCH

ŀ	<u>ئ</u> ئ	/ <u>}</u>	Ç Y	Ċ	33	٠.	3 2	and the same
		_	_					

Feedback Report a problem Satisfaction survey

Terms used attribute cache and TLB

Found **690** of **171,143** 

next

Sort results by Display

results

relevance -

expanded form

Save results to a Binder Search Tips

Open results in a new

Try an Advanced Search Try this search in The ACM Guide

window

Results 21 - 40 of 200

Result page: previous 1 2 3 4 5 6 7 8 9 10

Best 200 shown

Relevance scale 🔲 📟 🖬

21 Energy efficient memory systems: Generating physical addresses directly for saving instruction TLB energy

I. Kadayif, A. Sivasubramaniam, M. Kandemir, G. Kandiraju, G. Chen

November 2002 Proceedings of the 35th annual ACM/IEEE international symposium on Microarchitecture

**Publisher: IEEE Computer Society Press** 

Full text available: pdf(1.26 MB) Additional Information: full citation, abstract, references, citings, index Publisher Site

Power consumption and power density for the Translation Lookaside Buffer (TLB) are important considerations not only in its design, but can have a consequence on cache design as well. This paper embarks on a new philosophy for reducing the number of accesses to the instruction TLB (iTLB) for power and performance optimizations. The overall idea is to keep a translation currently being used in a register and avoid going to the iTLB as far as possible --- until there is a page change. We propose f ...

22 Poster Session 3: TLB and snoop energy-reduction using virtual caches in low-power \

chip-multiprocessors

Magnus Ekman, Per Stenström, Fredrik Dahlgren

August 2002 Proceedings of the 2002 international symposium on Low power electronics and design

**Publisher: ACM Press** 

Full text available: pdf(84.59 KB)

Additional Information: full citation, abstract, references, citings, index terms

In our quest to bring down the power consumption in low-power chip-multiprocessors, we have found that TLB and snoop accesses account for about 40% of the energy wasted by all L1 data-cache accesses. We have investigated the prospects of using virtual caches to bring down the number of TLB accesses. A key observaction is that while the energy wasted in the TLBs are cut, the energy associated with snoop accesses becomes higher. We then contribiute with two techniques to reduce the number of snoop ...

Keywords: CMP, low-power, snoop, virtual caches

23 MEDEA workshop: A transparent Linux super page kernel for Alpha, Sparc64 and

▲ IA32: reducing TLB misses of applications

Naohiko Shimizu, Ken Takatori

March 2003 ACM SIGARCH Computer Architecture News, Volume 31 Issue 1

**Publisher: ACM Press** 

Full text available: pdf(980.41 KB) Additional Information: full citation, abstract, references, index terms

Modern processors have various features for latency tolerance such as Hit-under-miss, Out-of-order execution, or Multi-threading. However, many processors must make a precise trap for a TLB miss, because they maintain the TLB with software and cannot distinguish the TLB scarcity from the page fault. It is very important for the application and/or the operating system to avoid the TLB misses as much as possible. Many processors have some super page features that extend the coverage of the TLB sig ...

24 Energy-aware compiling and scheduling: Compiler-directed code restructuring for

e reducing data TLB energy

M. Kandemir, I. Kadayif, G. Chen

September 2004 Proceedings of the 2nd IEEE/ACM/IFIP international conference on Hardware/software codesign and system synthesis

**Publisher: ACM Press** 

Full text available: pdf(172.79 KB) Additional Information: full citation, abstract, references, index terms

Prior work on TLB power optimization considered circuit and architectural techniques. A recent software-based technique for data TLBs has considered the possibility of storing the frequently used virtual-to-physical address translations in a set of translation registers (TRs), and using them when necessary instead of going to the data TLB. This paper presents a compiler-based strategy for increasing the effectiveness of TRs. The idea is to restructure the application code in such a fashion that ...

Keywords: code restructuring

25 Adapting Radix Sort to the Memory Hierarchy

Naila Rahman, Rajeev Raman

December 2001 Journal of Experimental Algorithmics (JEA), Volume 6

**Publisher: ACM Press** 

Full text available: pdf(446.81 KB)

ps(360.14 KB) tar(706.56 KB)

Additional Information: full citation, abstract, references, citings, index

terms

We demonstrate the importance of reducing misses in the translation-lookaside buffer (TLB) for obtaining good performance on modern computer architectures. We focus on least-significantbit first (LSB) radix sort, standard implementations of which make many TLB misses. We give three techniques which simultaneously reduce cache and TLB misses for LSB radix sort: reducing working set size, explicit block transfer and pre-sorting. We note that: • All the techniques above yield algorithms whose ...

**Keywords**: cache, efficient sorting algorithms, external-memory algorithms, locality of reference, memory hierarchy, radix sort, translation-lookaside buffer (TLB)

26 Translation lookaside buffer consistency: a software approach

D. L. Black, R. F. Rashid, D. B. Golub, C. R. Hill

April 1989 ACM SIGARCH Computer Architecture News, Proceedings of the third international conference on Architectural support for programming languages and operating systems ASPLOS-III, Volume 17 Issue 2

**Publisher: ACM Press** 

Full text available: pdf(1.38 MB)

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> terms

We discuss the translation lookaside buffer (TLB) consistency problem for multiprocessors, and introduce the Mach shootdown algorithm for maintaining TLB consistency in software. This algorithm has been implemented on several multiprocessors, and is in regular production use. Performance evaluations establish the basic costs of the algorithm and show that it has minimal impact on application performance. As a result, TLB consistency

does not pose an insurmountable obstacle to multiprocessor ... 27 Characterizing the memory behavior of Java workloads: a structured view and opportunities for optimizations Yefim Shuf, Mauricio J. Serrano, Manish Gupta, Jaswinder Pal Singh June 2001 ACM SIGMETRICS Performance Evaluation Review, Proceedings of the 2001 ACM SIGMETRICS international conference on Measurement and modeling of computer systems SIGMETRICS '01, Volume 29 Issue 1 **Publisher: ACM Press** Full text available: pdf(1.55 MB) Additional Information: full citation, abstract, references, citings This paper studies the memory behavior of important Java workloads used in benchmarking Java Virtual Machines (JVMs), based on instrumentation of both application and library code in a state-of-the-art JVM, and provides structured information about these workloads to help quide systems' design. We begin by characterizing the inherent memory behavior of the benchmarks, such as information on the breakup of heap accesses among different categories and on the hotness of references to fields and met ... 28 Options for dynamic address translation in COMAs Xiaogang Qiu, Michel Dubois April 1998 ACM SIGARCH Computer Architecture News , Proceedings of the 25th annual international symposium on Computer architecture ISCA '98, Volume Publisher: IEEE Computer Society, ACM Press Full text available: pdf(1.37 MB) Additional Information: full citation, abstract, references, citings, index Publisher Site In modern processors, the dynamic translation of virtual addresses to support virtual memory is done before or in parallel with the first-level cache access. As processor technology improves at a rapid pace and the working sets of new applications grow insatiably the latency and bandwidth demands on the TLB (Translation Lookaside Buffer) are getting more and more difficult to meet. The situation is worse in multiprocessor systems, which run larger applications and are plagued by the TLB consiste ... 29 High-bandwidth address translation for multiple-issue processors Todd M. Austin, Gurindar S. Sohi May 1996 ACM SIGARCH Computer Architecture News , Proceedings of the 23rd annual international symposium on Computer architecture ISCA '96, Volume Publisher: ACM Press Additional Information: full citation, abstract, references, citings, index Full text available: pdf(1.56 MB) terms In an effort to push the envelope of system performance, microprocessor designs are continually exploiting higher levels of instruction-level parallelism, resulting in increasing bandwidth demands on the address translation mechanism. Most current microprocessor designs meet this demand with a multi-ported TLB. While this design provides an excellent hit rate at each port, its access latency and area grow very quickly as the number of ports is increased. As bandwidth demands continue to increase ... 30 Modifying VM hardware to reduce address pin requirements Matthew Farrens, Arvin Park, Gary Tyson December 1992 ACM SIGMICRO Newsletter, Proceedings of the 25th annual international symposium on Microarchitecture MICRO 25, Volume 23 Issue Publisher: IEEE Computer Society Press, ACM Press

Full text available: 📆 pdf(607.69 KB) Additional Information: full citation, references, citings, index terms

31	Tradeoffs in supporting two page sizes	
٨	Madhusudhan Talluri, Shing Kong, Mark D. Hill, David A. Patterson	
~	April 1992 ACM SIGARCH Computer Architecture News, Proceedings of the 19th	
	annual international symposium on Computer architecture ISCA '92, Volume	
	20 Issue 2 Publisher: ACM Press	
	A delitional information full station whether the formation industrial	
	Full text available: pdf(1.18 MB)  Additional information: full citation, abstract, references, citings, index terms	
	As computer system main memories get larger and processor cycles-per-instruction (CPIs) get smaller, the time spent in handling translation lookaside buffer (TLB) misses could become a performance bottleneck. We explore relieving this bottleneck by (a) increasing the page size and (b) supporting two page sizes. We discuss how to build a TLB to support two page sizes and examine both alternatives experimentally with a dozen uniprogrammed, user-mode traces for the SPARC architectur	
32	Architectural support for translation table management in large address space	
	<u>machines</u>	
~	Jerry Huck, Jim Hays	
	May 1993 ACM SIGARCH Computer Architecture News, Proceedings of the 20th	
	annual international symposium on Computer architecture ISCA '93, Volume	
	21 Issue 2 Publisher: ACM Press	
	Full text available: pdf(1.34 MB)  Additional Information: full citation, abstract, references, citings, index terms	
	Virtual memory page translation tables provide mappings from virtual to physical	
	addresses. When the hardware controlled Translation Lookaside Buffers (TLBs) do not contain a translation, these tables provide the translation. Approaches to the structure and management of these tables vary from full hardware implementations to complete software based algorithms. The size of the virtual address space used by processes is rapidly growing beyond 32 bits of address. As the utilized ad	
33	Docian tradeoffs for software managed TLPs	
•	Design tradeoffs for software-managed TLBs  David Nagle, Richard Uhlig, Tim Stanley, Stuart Sechrest, Trevor Mudge, Richard Brown  May 1993 ACM SIGARCH Computer Architecture News, Proceedings of the 20th  annual international symposium on Computer architecture ISCA '93, Volume  21 Issue 2  Publisher: ACM Press	
	Full text available: pdf(1.14 MB)  Additional Information: full citation, abstract, references, citings, index terms	
	An increasing number of architectures provide virtual memory support through software-managed TLBs. However, software management can impose considerable penalties, which are highly dependent on the operating system's structure and its use of virtual memory. This work explores software-managed TLB design tradeoffs and their interaction with a range of operating systems including monolithic and microkernel designs. Through hardware monitoring and simulations, we explore TLB performance for be	
34 <b>②</b>	A new page table for 64-bit address spaces  M. Talluri, M. D. Hill, Y. A. Khalidi  December 1995 ACM SIGOPS Operating Systems Review, Proceedings of the fifteenth  ACM symposium on Operating systems principles SOSP '95, Volume 29	
	Issue 5	
	Publisher: ACM Press	
	Full text available: pdf(1.97 MB) Additional Information: full citation, references, citings, index terms	
25		
35	Optimal allocation of on-chip memory for multiple-API operating systems  D. Nagle, R. Uhlig, T. Mudge, S. Sechrest	

April 1994 ACM SIGARCH Computer Architecture News, Proceedings of the 21ST annual international symposium on Computer architecture ISCA '94, Volume 22 Issue 2

Publisher: IEEE Computer Society Press, ACM Press

Full text available: pdf(1.27 MB)

Additional Information: full citation, abstract, references, citings, index terms

The allocation of die area to different processor components is a central issue in the design of single-chip microprocessors. Chip area is occupied by both core execution logic, such as ALU and FPU datapaths, and memory structures, such as caches, TLBs, and write buffers. This work focuses on the allocation of die area to memory structures through a cost/benefit analysis. The cost of memory structures with different sizes and associativities is estimated by using an established area model for on ...

36 Optimizing database architecture for the new bottleneck: memory access

Stefan Manegold, Peter A. Boncz, Martin L. Kersten

December 2000 The VLDB Journal — The International Journal on Very Large Data

Bases, Volume 9 Issue 3

Publisher: Springer-Verlag New York, Inc.

Full text available: pdf(357.33 KB) Additional Information: full citation, abstract, citings, index terms

In the past decade, advances in the speed of commodity CPUs have far out-paced advances in memory latency. Main-memory access is therefore increasingly a performance bottleneck for many computer applications, including database systems. In this article, we use a simple scan test to show the severe impact of this bottleneck. The insights gained are translated into guidelines for database architecture, in terms of both data structures and algorithms. We discuss how vertically fragmented data struc ...

Keywords: Decomposed storage model, Implementation techniques, Join algorithms, Main-memory databases, Memory access optimization, Query processing

37 Cache Memories



Alan Jay Smith

September 1982 ACM Computing Surveys (CSUR), Volume 14 Issue 3

Publisher: ACM Press

Full text available: pdf(4.61 MB)

Additional Information: full citation, references, citings, index terms

38 The use of multithreading for exception handling

Craig B. Zilles, Joel S. Emer, Gurindar S. Sohi

November 1999 Proceedings of the 32nd annual ACM/IEEE international symposium on Microarchitecture

**Publisher: IEEE Computer Society** 

Full text available: pdf(1.49 MB) Additional Information: full citation, abstract, references, citings, index terms Publisher Site

Common hardware exceptions, when implemented by trapping, unnecessarily serialize program execution in dynamically scheduled superscalar processors. To avoid the consequences of trapping the main program thread, multithreaded CPUs can exploit control and data independence by executing the exception handler in a separate hardware context. The main thread doesn't squash instructions after the excepting instruction, conserving fetch bandwidth and allowing execution of instructions inde ...

39 Design tradeoffs for software-managed TLBs



Richard Uhlig, David Nagle, Tim Stanley, Trevor Mudge, Stuart Sechrest, Richard Brown August 1994 ACM Transactions on Computer Systems (TOCS), Volume 12 Issue 3

Publisher: ACM Press

Additional Information:

Full text available: pdf(1.85 MB)

full citation, abstract, references, citings, index terms, review

An increasing number of architectures provide virtual memory support through software-managed TLBs. However, software management can impose considerable penalties that are highly dependent on the operating system's structure and its use of virtual memory. This work explores software-managed TLB design tradeoffs and their interaction with a range of monolithic and microkernel operating systems. Through hardware monitoring and simulation, we explore TLB performance for benchmarks running on a ...

**Keywords**: hardware monitoring, translation lookaside buffer (TLB), trap-driven simulation

40	t of node size on the performance of cache-conscious B±-trees							
	Richard A. Hankins, Jignesh M. Patel							
4	Richard A. Hankins, Jignesh M. Patel June 2003 ACM SIGMETRICS Performance Evaluation Review, Proceedings of the							
	2003 ACM SIGMETRICS international conference on Measurement a							
	modeling of computer systems SIGMETRICS '03, Volume 31 Issue 1							

Publisher: ACM Press

Full text available: pdf(271.16 KB) Additional Information: full citation, abstract, references, index terms

In main-memory databases, the number of processor cache misses has a critical impact on the performance of the system. Cache-conscious indices are designed to improve performance by reducing the number of processor cache misses that are incurred during a search operation. Conventional wisdom suggests that the index's node size should be equal to the cache line size in order to minimize the number of cache misses and improve performance. As we show in this paper, this design choice ignores additi ...

**Keywords**: B+-tree, cache-conscious, index

Results 21 - 40 of 200 Result page: previous 1 2 3 4 5 6 7 8 9 10 next

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.

Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Player